

09-05-00

A

Please type a plus sign (+) inside this box → ☐PTO/SB/05 (4/98)  
Approved for use through 09/30/2000. OMB 0651-0032  
Patent and Trademark Office: U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number

**UTILITY  
PATENT APPLICATION  
TRANSMITTAL**

(Only for new nonprovisional applications under 37 C.F.R. § 1.53(b))

Attorney Docket No. CISCO-2357

First Inventor or Application Identifier Gourlay et al.

Title Phased Learning Approach to determining...

Express Mail Label No. EL575423338US

**APPLICATION ELEMENTS**

See MPEP chapter 600 concerning utility patent application contents.

1. ☐ \* Fee Transmittal Form (e.g., PTO/SB/17)  
(Submit an original and a duplicate for fee processing)
2. ☒ Specification [Total Pages 32]
  - Descriptive title of the Invention
  - Cross References to Related Applications
  - Statement Regarding Fed sponsored R & D
  - Reference to Microfiche Appendix
  - Background of the Invention
  - Brief Summary of the Invention
  - Brief Description of the Drawings (if filed)
  - Detailed Description
  - Claim(s)
  - Abstract of the Disclosure
3. ☒ Drawing(s) (35 U.S.C. 113) [Total Sheets 4]
  - a. ☐ Newly executed (original or copy)
  - b. ☐ Copy from a prior application (37 C.F.R. § 1.63(d))  
(for continuation/divisional with Box 16 completed)
  - i. ☐ **DELETION OF INVENTOR(S)**  
Signed statement attached deleting inventor(s) named in the prior application, see 37 C.F.R. §§ 1.63(d)(2) and 1.33(b).
4. Oath or Declaration [Total Pages ]
  - a. ☐ Newly executed (original or copy)
  - b. ☐ Copy from a prior application (37 C.F.R. § 1.63(d))  
(for continuation/divisional with Box 16 completed)
  - i. ☐ **DELETION OF INVENTOR(S)**  
Signed statement attached deleting inventor(s) named in the prior application, see 37 C.F.R. §§ 1.63(d)(2) and 1.33(b).

\* NOTE FOR ITEMS 1 & 13: IN ORDER TO BE ENTITLED TO PAY SMALL ENTITY FEES, A SMALL ENTITY STATEMENT IS REQUIRED (37 C.F.R. § 1.27), EXCEPT IF ONE FILED IN A PRIOR APPLICATION IS RELIED UPON (37 C.F.R. § 1.28).

**ADDRESS TO:**Assistant Commissioner for Patents  
Box Patent Application  
Washington, DC 20231

5. ☐ Microfiche Computer Program (Appendix)
6. Nucleotide and/or Amino Acid Sequence Submission (if applicable, all necessary)
  - a. ☐ Computer Readable Copy
  - b. ☐ Paper Copy (identical to computer copy)
  - c. ☐ Statement verifying identity of above copies

**ACCOMPANYING APPLICATION PARTS**

7. ☐ Assignment Papers (cover sheet & document(s))
8. ☐ 37 C.F.R. § 3.73(b) Statement ☐ Power of Attorney  
(when there is an assignee)
9. ☐ English Translation Document (if applicable)
10. ☐ Information Disclosure Statement (IDS)/PTO-1449 ☐ Copies of IDS Citations
11. ☐ Preliminary Amendment
12. ☒ Return Receipt Postcard (MPEP 503)  
(Should be specifically itemized)
13. ☐ \* Small Entity Statement filed in prior application, Status still proper and desired  
(PTO/SB/09-12)
14. ☐ Certified Copy of Priority Document(s)  
(if foreign priority is claimed)
15. ☐ Other: \_\_\_\_\_

16. If a CONTINUING APPLICATION, check appropriate box, and supply the requisite information below and in a preliminary amendment:

☐ Continuation ☐ Divisional ☐ Continuation-in-part (CIP) of prior application No: \_\_\_\_\_

Prior application information: Examiner \_\_\_\_\_ Group / Art Unit: \_\_\_\_\_

For CONTINUATION or DIVISIONAL APPS only: The entire disclosure of the prior application, from which an oath or declaration is supplied under Box 4b, is considered a part of the disclosure of the accompanying continuation or divisional application and is hereby incorporated by reference. The incorporation can only be relied upon when a portion has been inadvertently omitted from the submitted application parts.

**17. CORRESPONDENCE ADDRESS**☐ Customer Number or Bar Code Label

(Insert Customer No. or Attach bar code label here)

or ☐ Correspondence address below

Name	Marc S. Hanish				
	D'Alessandro & Ritchie				
Address	P.O. Box 640640				
City	San Jose	State	CA	Zip Code	95164-0640
Country	USA	Telephone	408-441-1100	Fax	408-441-8400

Name (Print/Type)	Marc S. Hanish	Registration No. (Attorney/Agent)	42,262
Signature	<i>Marc S. Hanish</i>	Date	8/31/00

Burden Hour Statement: This form is estimated to take 0.2 hours to complete. Time will vary depending upon the needs of the individual case. Any comments on the amount of time you are required to complete this form should be sent to the Chief Information Officer, Patent and Trademark Office, Washington, DC 20231. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Assistant Commissioner for Patents, Box Patent Application, Washington, DC 20231.

08/31/00

1c821 U.S. PTO

09/05/00 09/05/00

1c821 U.S. PTO  
09/05/00

08/31/00

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Application of:

Gourlay et al.

Serial No.: [Not yet assigned]

Filed: August 31, 2000

For: PHASED LEARNING APPROACH  
TO DETERMINING CLOSEST  
CONTENT SERVING SITES

Art Unit: Not yet assigned

Examiner: Not yet assigned



## CERTIFICATE OF MAILING

"Express Mail" mailing label no: EL575423338US

Date of Deposit: August 31, 2000

I hereby certify that this correspondence is being deposited  
with the United States Postal Service "Express Mail Post  
Office to Addressee" service under 37 CFR 1.10 on the  
date indicated above and is addressed to:

Box Patent Application  
Assistant Commissioner for Patents  
Washington, D.C. 20231

*Lola R. Salazar*  
Lola R. Salazar

**TRANSMITTAL LETTER**

Honorable Assistant Commissioner  
for Patents  
Box Patent Application  
Washington, D.C. 20231

Dear Sir:

Enclosed for filing please find the patent application for an invention entitled,  
"PHASED LEARNING APPROACH TO DETERMINING CLOSEST CONTENT  
SERVING SITES", filed on behalf of Cisco Technologies, Inc., assignee from  
inventor(s) Douglas Gourlay and James Aviani, including Utility Patent Application

Transmittal, 20 pages of specification, 11 pages of claims, 4 sheets of drawing figures,  
and 1 page of Abstract.

The attorney's Docket Number is CISCO-2357

Kindly address all communications regarding this application to:

Marc S. Hanish  
D'Alessandro & Ritchie  
P.O. Box 640640  
San Jose, CA 95164-0640  
Telephone (408) 441-1100

No fee is being paid at this time.

In the event any variance exists between the amount enclosed and the Patent  
Office charges for filing the above-noted documents, the Assistant Commissioner is  
hereby authorized to charge or credit the difference to our Deposit Account No. 04-0025.  
A duplicate of this page is enclosed.

Respectfully submitted,  
D'ALESSANDRO & RITCHIE



Marc S. Hanish  
Reg. No. 42,262

Dated: August 31, 2000

D'Alessandro & Ritchie  
P.O. Box 640640  
San Jose, CA 95164-0640  
(408) 441-1100

This application is submitted in the name of inventors Douglas Gourlay and James Aviani, both assignors to Cisco Technology, Inc., a California Corporation.

5

## S P E C I F I C A T I O N

### **PHASED LEARNING APPROACH TO DETERMINING CLOSEST CONTENT SERVING SITES**

10

#### BACKGROUND OF THE INVENTION

##### Field of the Invention

The present invention relates to the field of network communications. More specifically, the present invention relates to the use of a phased learning approach to determine the closest content serving site to a user or server.

15

##### The Background

In a distributed computer network, content sites may be distributed throughout a network, each site containing identical content to at least one of the other distributed sites. This allows for both load balancing among multiple sites as well as for faster connections as users may connect to a server which is closest (in terms of transit time) to their own computer.

20

This design has gained acceptance on the Internet, where it is routine for large content distributors to maintain multiple servering locations throughout the Internet.

25

When one server is busy, traffic may merely be directed to a different server. Likewise, if one server is down for repairs, traffic may merely be directed to a different server.

Finally, and perhaps most importantly, when multiple servers are available, the traffic may merely be directed to the server that is closest to the user.

When a user requests content that is located on multiple distributed servers, the server having the lowest round-trip response time to the user's DNS Server must be determined. FIG. 1 is a block diagram illustrating a situation where a user requests content located on multiple distributed servers. PC 2 may be attempting to access content at "www.example.com". Thus, it first requests Domain Naming System (DNS) resolution from Client DNS Server 4. Then Client DNS Server 4 queries the root DNS to determine the authoritative DNS server for example.com.

The Client DNS Server 4 may then query the authoritative DNS server for the domain www.example.com, which returns an NS record for www.example.com. The Client DNS server 4 may then query the global server load balancer 6 for the "www.example.com" domain. At the global server load balancer 6, the various servers 8a, 8b, 8c containing the content are known. The global server load balancer 6 then may query each of the servers 8a, 8b, 8c. The sites then respond by each sending a TCP/IP syn/ack packet to the client DNS server 4. The client DNS server 4 then receives the TCP/IP syn/ack packet and sends a reset (RST) packet back to each server 8a, 8b, 8c, resetting the syn/ack. This flows from the Client DNS server 4 to each server 8a, 8b, 8c directly.

When the servers 8a, 8b, 8c each receive the RST packet back, they may compute their respective transit times (by using the difference between the time they originally sent the syn/ack packet and the time they received the RST packet back). This information may then be passed to the global server load balancer 6. The global server

5 load balancer may then compare the transit times of all of the servers 8a, 8b, 8c, determine which is the lowest, and inform the client DNS server 4 of the resolved IP address of the lowest timed site. The client DNS server 4 may create a record to the PC containing the information indicating the server with the lowest transit time, at which point the PC may connect to the IP address for that server to retrieve the content.

This can be a fairly time consuming process, especially when the number of servers hosting the content grows larger. The relative lack of speed of the process may be bearable when there are only 3 servers hosting the content, but when there are 3000, the delays while waiting for servers to respond can be costly. Thus, what is needed is a

15 solution that directs a user to the appropriate server in the least amount of time, especially as the number of servers grows large.

SUMMARY OF THE INVENTION

In order to direct content requests to an appropriate content serving site in a computer network, a phased learning approach is utilized to ensure that, as best as possible, the request is made to content serving site with the shortest delay. In a setup phase, an indirect path return geographic sever load balancer times queries to all of the individual content serving sites so that the queries all arrive at the content serving sites at the same time. Therefore, when the requesting fixed location receives communications from the content serving sites, it can easily tell which content serving site has the least delay by an established race condition. The winner of the race may then be relayed to the indirect path return geographic server load balancer for later usage. In an execution mode, only the  $m$  fastest content serving sites and  $n$  other sites (used to test random and new sites) are queried when a DNS request arrives from the requesting fixed location. The particular  $m$  fastest content serving sites and  $n$  other sites may be dynamically updated so as to ensure the most reliable directing of requests. This solution provides a very efficient and effective means by which to determine closest content serving sites while keeping load balancer-created traffic at a minimum.

### BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a block diagram illustrating a situation where a user requests content located on multiple distributed servers.

5        FIG. 2 is a block diagram illustrating a system utilizing indirect path return geographic server load balancing technology in accordance with a specific preferred embodiment of the present invention.

10       FIG. 3 is a flow diagram illustrating a method for using a phased learning approach for determining the closest content serving sites to a fixed location in a computer network in accordance with a specific embodiment of the present invention.

15       FIG. 4 is a block diagram illustrating an apparatus for using a phased learning approach for determining the closest content serving sites to a fixed location in a computer network in accordance with a specific embodiment of the present invention.



DETAILED DESCRIPTION OF A PREFERRED EMBODIMENT

In the following description, a preferred embodiment of the invention is described with regard to preferred process steps and data structures. However, those skilled in the art will recognize, after perusal of this application, that embodiments of the invention may be implemented using at least one general purpose computer operating under program control and/or using a program storage device, and that modification of the general purpose computer to implement the components, process steps, and/or data structures described herein would not require undue invention.

In accordance with a specific embodiment of the present invention, the components, process steps, and/or data structures are implemented using software. This implementation is not intended to be limiting in any way. Different implementations may be used and may include other types of operating systems, computing platforms, and/or computer programs. In addition, those of ordinary skill in the art will readily recognize that devices of a less general purpose nature, such as hardwired devices, devices relying on FPGA (field programmable gate array) or ASIC (application specific integrated circuit) technology, or the like, may also be used without departing from the scope and spirit of the inventive concepts disclosed herewith.

The present invention utilizes a phased learning approach along with indirect path return geographic server load balancing technology in order to quickly and efficiently provide a user with a list of the fastest servers hosting the requested content.

5           FIG. 2 is a block diagram illustrating a system utilizing indirect path return geographic server load balancing technology in accordance with a specific embodiment of the present invention. PC 50 attempts to access content at "www.example.com". Thus, it first requests Domain Naming System (DNS) resolution from Client DNS Server 52. Then Client DNS Server 52 queries the root DNS to determine the authoritative DNS  
10           Server for the requested domain. The Client DNS Server 52 may then query the authoritative DNS server for the domain "www.example.com". The Client DNS server 52 may then query the indirect path return geographic server load balancer 54 for the "www.example.com" domain. At the indirect path return geographic server load balancer 54, the various servers 56a, 56b, 56c containing the content are known. Additionally, the  
15           time it takes to send a packet from the indirect path return geographic server load balancer 54 to each of the various servers 56a, 56b, 56c is also known. This may be determined by pinging each of the servers and measuring the response times (the time it takes to receive a response back). Pinging may be done at regular intervals, to have on  
20           hand the most up-to-date information as to the response times of the various servers 56a, 56b, 56c.

With the knowledge of the response times of each server 56a, 56b, 56c, the indirect path return geographic server load balancer 54 may then precisely time queries to each of the servers 56a, 56b, 56c so that the queries are received at each of the servers 56a, 56b, 56c at exactly the same time. Thus if pinging determines that the response time for the servers 56a, 56b, and 56c are 100ms, 20ms, and 50ms, respectively, then the indirect path return geographic server load balancer 54 will first send a query to server 56a, then wait 25ms and send a query to server 56c (the difference in the response times between server 56a and 56c divided by 2 because the query is only traveling one-way), then wait another 15ms and send a query to server 56b (difference in the response times between server 56a and 56b is 80ms, divided by 2 gives 40ms from the time the query to server 56a was sent). This ensures that the queries arrive at each of the servers 56a, 56b, and 56c at the same time.

Since the queries arrive at each of the servers 56a, 56b, 56c at the same time, the DNS responses from 56a, 56b, 56c to the client DNS server 52 may be sent at the same time. This creates what may be called a one-way race condition. In order to determine the server 56a, 56b, 56c with the fastest transit time, the client DNS server 52 need only determine the server from which the first DNS Response packet it receives came. When a server 56a, 56b, or 56c wins the race it is known to the server by the corresponding HTTP Request coming in from an IP address with similarity to the Client-DNS server. The winning of the one-way race triggers an update to the indirect path geographic server load balancer with the client DNS servers IP address and the transit time for the race to

occur. The indirect path geographic server load balancer then adds this to an entry in its content routing table where it keeps the client-DNS server IP address mapped to the *m* fastest sites that responded to it.

5 One problem, however, with this solution occurs when a large number of servers are placed in the system. The indirect path return geographic server load balancer 54 must send requests to many or all of them, thus creating traffic that may actually interfere with the efficiency of the system, especially the client-DNS server. So while the time it takes to determine which is the fastest server when a user wishes to access content is  
10 greatly decreased, the overall speed of the servers may actually be slower.

In order to remedy this concern, only a certain number of servers will be "active" at any one time for each client-DNS server, treated as independent entities. "Active" means that the server is subject to receive a query from the indirect path return  
15 geographic server load balancer 54 for a particular client DNS server. In a specific embodiment of the present invention, 5 servers will be active at any one time for each client DNS server, although there can be overlap between different client DNS servers. The 5 servers will made up of the 3 servers having the lowest response time to the client-DNS server as well as 2 other servers used for testing (which will be described below).

20 Determining which servers are active will be performed using a phased learning approach. When the indirect path return geographic server load balancer 54 is first

connected to the network or the network is otherwise "set up", the system is deemed to be in the setup phase. During the setup phase, all of the servers are sequentially tested as traffic comes in from different client-DNS IP addresses. All of the servers are pinged from the indirect path geographic server load balancer to determine their relative distance

5 from the indirect path return geographic server load balancer. This data is stored in records at the indirect path return geographic server load balancer 54, each record comprising the halved ping time from a content server to the indirect path return geographic server load balancer. As the indirect path return geographic server load balancer receives the first request from a client-DNS server, it randomly selects 10

10 content servers to forward the equalized request to. These 10 servers respond back to the client DNS server establishing a race. The winner of this race is reported back to the indirect path return geographic server load balancer by the content server. The indirect path return geographic server load balancer 54 then adds a record to its table with the winning content server as an entry under the client DNS servers IP address with the

15 transit time for it to respond and receive traffic. If space is a concern, then the number of records may be limited such as storing only the 3 fastest content servers and simply replacing a record corresponding to a particular content server if a faster server is found.

After the system has been set up, it still may be necessary to adjust which servers

20 are active based on changes that occur within the network. Typical changes include high traffic to a particular server and addition or deletion of a server from the network. When the system is not in the setup phase, it is in the execution phase. During the execution

phase, the 3 fastest servers are active at all times for each client-DNS server. The two extra slots are used for testing random servers in a stable network (no servers added or deleted). If one of the random servers tested winds up having a lower transit time than one of the three fastest servers, then it will supplant the slowest of the three fastest

5 servers. If a new server is added to the system, it will be immediately added to one of the two extra slots so it can be tested until has become one of the fastest servers mapped to a client-DNS server. Just like with the random servers, if the new server tests out with a lower transit time than one of the three fastest servers, then it will supplant the slowest of the three fastest servers.

10 It should be noted that it is not necessary that the number of active servers be 5. Any number will work, and the number may have to be varied for different applications in different systems. Additionally, the number of extra slots to use may also be varied.

15 Furthermore, it may be advantageous at times to make the number of active servers dynamic. If there is a high rate of change in the list of fastest servers, it indicates that the network may not be stable. This may be used as an indicator to a system administrator that action needs to be taken to examine the root of the instability in the system. During this instability, it may be preferable to increase the number of servers in

20 the list of fastest servers so the rate of change can be mitigated as the available pool of servers to be tested for each iteration is increased. The number may be scaled back down when stability returns to the system.

FIG. 3 is a flow diagram illustrating a method for using a phased learning approach for determining the closest content serving sites to a fixed location in a computer network in accordance with a specific embodiment of the present invention.

- 5 This method may be performed by an indirect path return geographic server load balancer.

The system begins in a setup phase, when the computer network (or a portion thereof) is first setup. However, setup mode may also be entered into whenever an administrator wishes to test the network. Thus, if a new version of the indirect path return geographic server load balancer or the table has been erased, the administrator may wish to enter setup mode even though it is not the first time the network is setup.

The fixed location will generally be a PC and client DNS server combination. However, as the computer network grows more complex, there may be a need to serve content to other devices, such as gateways, routers, etc. Thus the term "fixed location" should not be limited to PCs and client DNS servers. At 100, a DNS request is received from the fixed location. Even though the system is in the setup phase, there still needs to be a fixed location defined for testing purposes. At 102, a response time is determined for each of the content serving sites. This is determined by pinging each of the content serving sites and then measuring the time it takes to receive a reply back from each of the content serving sites. That time is the response time for each content serving site.

At 104, queries are sent to each of the content serving sites, timing the queries so that they arrive at each of the content serving sites at the same time by using said response time for each of the content serving sites. This is accomplished by querying the content serving sites in order of their response times, longest response time first. The query for a particular content serving site is delayed from the query for the content serving site with the longest response time for a time equal to one half of the particular content serving site's response time.

At 106, data is received from the fixed location as to the transit times of each of the content serving sites. The fixed location would have measured these transit times by tracking when a response from each of the content serving sites was received (the response from each of the content serving sites being generated upon receipt of the query from the indirect path return geographic server load balancer). The data may be in a form of a record listing the IP address of fixed location, the transit time from a particular content serving site, and the IP address of that particular content serving site. At 108, the data may then be stored in table for later use.

When the network is operating normally, the system may be in the execution phase. At 110, a DNS request is received from the fixed location. At 112, a response time is determined for each of  $n$  fastest content serving sites and  $m$  other content serving sites, said  $n$  fastest content serving sites chosen by using said data and choosing the  $n$



content serving sites with the lowest transit times, said  $m$  other content serving sites chosen by selecting new content serving sites as well as randomly selecting old content serving sites. Thus, whenever a new content serving site is added to the network it will become one of the  $m$  other content serving sites until it has been tested once.

- 5 Determining the response time is accomplished by pinging each of the  $n$  fastest content serving sites and  $m$  other content serving sites and then measuring the time it takes to receive a reply back from each of the  $n$  fastest content serving sites and  $m$  other content serving sites. That time is the response time for each of  $n$  fastest content serving site and  $m$  other content serving sites.

At 114, queries are sent to each of  $n$  fastest content serving sites and  $m$  other content serving sites, timing the queries so that they arrive at each of the  $n$  fastest content serving sites and  $m$  other content serving sites at the same time by using said response time for each of the  $n$  fastest content serving sites and  $m$  other content serving sites. This is accomplished by querying the  $n$  fastest content serving sites and  $m$  other content serving sites in order of their response times, longest response time first. The query for a particular  $n$  fastest content serving sites or  $m$  other content serving sites is delayed from the query for the  $n$  fastest content serving sites or  $m$  other content serving sites with the longest response time for a time equal to one half of the particular content serving site's response time.

At 116, data is received from the fixed location as to the transit times of each of the  $n$  fastest content serving sites and  $m$  other content serving sites. The fixed location would have measured these transit times by tracking when a response from each of the  $n$  fastest content serving sites and  $m$  other content serving sites was received (the response from the  $n$  fastest content serving sites and  $m$  other content serving sites being generated upon receipt of the query from the indirect path return geographic server load balancer). The data may be in a form of a record listing the IP address of fixed location, the transit time from a particular content serving site, and the IP address of that particular content serving site. At 118, the data may then be stored in table for later use.

In a specific embodiment of the present invention,  $n$  is 3 and  $m$  is 2. However, as discussed earlier, in some situations it is preferable to have these numbers be dynamic.

FIG. 4 is a block diagram illustrating an apparatus for using a phased learning approach for determining the closest content serving sites to a fixed location in a computer network in accordance with a specific embodiment of the present invention. An indirect path return geographic server load balancer.150 may be located in a network with content serving sites 152a, 152b, 152c, 152d and a fixed location 154.

The system begins in the setup phase, when the computer network (or a portion thereof) is first setup. However, setup mode may also be entered into whenever an administrator wishes to test the network. Thus, if a new version of the indirect path

return geographic server load balancer or the table has been erased, the administrator may wish to enter setup mode even though it is not the first time the network is setup.

The fixed location will generally be a PC and client DNS server combination.

- 5 However, as the computer network grows more complex, there may be a need to serve content to other devices, such as gateways, routers, etc. Thus the term "fixed location" should not be limited to PCs and client DNS servers. A DNS request receiver 156 coupled to the fixed location 154 receives a DNS request from the fixed location 154. Even though the system is in the setup phase, there still needs to be a fixed location defined for testing purposes. A response time determiner 158 coupled to said DNS request receiver 156 and to the content serving sites 152a, 152b, 152c, 152d determines a response time for each of the content serving sites 152a, 152b, 152c, 152d. This is determined by using a pinger 160 to ping each of the content serving sites 152a, 152b, 152c, 152d and then measuring the time it takes to receive a reply back at a response time receiver 162 from each of the content serving sites. That time is the response time for each content serving site.

- 20 A query sender 164 coupled to said response timer determiner 162 and to the content serving sites 152a, 152b, 152c, 152d sends queries to each of the content serving sites 152a, 152b, 152c, 152d, timing the queries so that they arrive at each of the content serving sites at the same time by using said response time for each of the content serving sites. This is accomplished by ordering the content serving sites in order of their

response times, longest response time first, using a response timer orderer 166, and then timing the queries with a query timer 168. The query for a particular content serving site is delayed from the query for the content serving site with the longest response time for a time equal to one half of the particular content serving site's response time.

5

A data receiver 170 coupled to the fixed location 154 receives data from the fixed location 154 as to the transit times of each of the content serving sites. The fixed location 154 would have measured these transit times by tracking when a response from each of the content serving sites 152a, 152b, 152c, 152d was received (the response from each of the content serving sites 152a, 152b, 152c, 152d being generated upon receipt of the query from the indirect path return geographic server load balancer 150). The data may be in a form of a record listing the IP address of fixed location, the transit time from a particular content serving site, and the IP address of that particular content serving site. A data storer 172 coupled to said data receiver 170 and to a memory 174 may store the data in a table in memory 174 for later use.

When the network is operating normally, the system may be in the execution phase. In order to distinguish between the two phases, a phase determiner 176 coupled to said response time determiner 158, said query sender 164, and said data receiver 170, is provided. In the execution phase, the DNS request receiver 156 receives a DNS request from the fixed location 154. The response time determiner 158 may then determine a response time for each of  $n$  fastest content serving sites and  $m$  other content serving sites,

said  $n$  fastest content serving sites chosen by using said data and choosing the  $n$  content serving sites with the lowest transit times, said  $m$  other content serving sites chosen by selecting new content serving sites as well as randomly selecting old content serving sites. This choosing of the fastest content serving sites may be accomplished by using an  $n$  fastest content serving site chooser 178, which may access the memory 174 to aid in its determination. Additionally, an  $m$  other content serving site chooser 180 may be provided to choose which new/random sites to have active. Thus, whenever a new content serving site is added to the network it will become one of the  $m$  other content serving sites until it has been tested once.

Determining the response time is accomplished by pinging each of the  $n$  fastest content serving sites and  $m$  other content serving sites using the pinger 160 and then measuring the time it takes to receive a reply back at said response timer receiver 162 from each of the  $n$  fastest content serving sites and  $m$  other content serving sites. That time is the response time for each of  $n$  fastest content serving site and  $m$  other content serving sites.

The query sender 164 sends queries to each of  $n$  fastest content serving sites and  $m$  other content serving sites, timing the queries so that they arrive at each of the  $n$  fastest content serving sites and  $m$  other content serving sites at the same time by using said response time for each of the  $n$  fastest content serving sites and  $m$  other content serving sites. This is accomplished by querying the  $n$  fastest content serving sites and  $m$  other

content serving sites in order of their response times, longest response time first, using said response time orderer. The query for a particular  $n$  fastest content serving sites or  $m$  other content serving sites is delayed from the query for the  $n$  fastest content serving sites or  $m$  other content serving sites with the longest response time for a time equal to one

5 half of the particular content serving site's response time.

The data receiver 170 receives data from the fixed location 154 as to the transit times of each of the  $n$  fastest content serving sites and  $m$  other content serving sites. The fixed location 154 would have measured these transit times by tracking when a response

10 from each of the  $n$  fastest content serving sites and  $m$  other content serving sites was received (the response from the  $n$  fastest content serving sites and  $m$  other content serving sites being generated upon receipt of the query from the indirect path return geographic server load balancer). The data may be in a form of a record listing the IP address of

15 fixed location, the transit time from a particular content serving site, and the IP address of that particular content serving site. A data storer 172 coupled to said data receiver 170 and to a memory 174 may store the data in a table in memory 174 for later use.

In a specific embodiment of the present invention,  $n$  is 3 and  $m$  is 2. However, as discussed earlier, in some situations it is preferable to have these numbers be dynamic.

20 While embodiments and applications of this invention have been shown and described, it would be apparent to those skilled in the art having the benefit of this

20

CLAIMS

What is claimed is:

5 / 1. A method for using a phased learning approach for determining the closest content serving sites to a fixed location in a computer network including:

in a setup phase:

storing data as to the transit times of each of the content serving sites in a table;

10 in an execution phase:

receiving a DNS request from said fixed location;

determining a response time for each of  $n$  fastest content serving sites and  $m$  other content serving sites, said  $n$  fastest content serving sites chosen by using said data and choosing the  $n$  content serving sites with the lowest transit times, said  $m$  other content serving sites chosen by selecting new content serving sites as well as randomly selecting old content serving sites;

15 sending queries to each of the  $n$  fastest content serving sites and  $m$  other content serving sites, timing said queries so that they arrive at each of the  $n$  fastest content serving sites and  $m$  other content serving sites at the same time by using said response time for each of the  $n$  fastest content serving sites and  $m$  other content serving sites;

20



receiving data from said fixed location as to the transit times of each of the  
 $n$  fastest content serving sites and  $m$  other content serving sites; and  
 storing said data in said table.

- 5    2.    The method of claim 1, wherein said setup phase includes:  
          receiving a DNS request from said fixed location;  
          determining a response time for each of the content serving sites;  
          sending queries to each of the content serving sites, timing said queries so that  
          they arrive at each of the content serving sites at the same time by using said response  
 10    time for each of the content serving sites;  
          receiving data from said fixed location as to the transit times of each of the content  
          serving sites;  
          storing said data in a table;
- 15    3.    The method of claim 2, wherein said determining the response time for each of the  
          content serving sites includes:  
          pinging each of the content serving sites;  
          receiving a reply from each of the content serving sites; and  
          determining the response time for each of the content serving sites by computing  
 20    the difference in time from said pinging to said receiving for each of the content serving  
          sites.

4. The method of claim 1, wherein said determining the response time for each of  $n$  fastest content serving sites and  $m$  other content serving sites includes:

pinging each of  $n$  fastest content serving sites and  $m$  other content serving sites;

receiving a reply from each of  $n$  fastest content serving sites and  $m$  other content

5 serving sites; and

determining the response time for each of  $n$  fastest content serving sites and  $m$  other content serving sites by computing the difference in time from said pinging to said receiving for each of the each of said  $n$  fastest content serving sites and  $m$  other content serving sites.

5. The method of claim 1, wherein transit time is the time it takes a content serving site to send a packet to said original location after receiving a query.

6. The method of claim 1, wherein said data includes at least one record, said record including a transit time, a network address for said fixed location, and a network address for a content serving site, said transit time being the transit time between said fixed location and said content serving site.

7. The method of claim 2, wherein said sending queries to each of the content serving sites includes:

querying each of the content serving sites in order of their response times, longest response time first, setting a delay for a query for a particular content serving site from

the time the query to the content serving site with the longest response time as one half of said particular content serving site's response time.

8. The method of claim 1, wherein said sending queries to each of the  $n$  fastest

5 content serving sites and  $m$  other content serving sites includes:

querying each of said content serving sites  $n$  fastest content serving sites and  $m$  other content serving sites in order of their response times, longest response time first, setting a delay for a query for a particular content serving site from the time the query to said  $n$  fastest content serving site or  $m$  other content serving site with the longest response  
10 time as one half of said particular content serving site's response time.

9. The method of claim 1, wherein  $n$  is 3 and  $m$  is 2.

10. The method of claim 1, wherein  $n$  may be varied based on stability of the  
15 computer network.

11. The method of claim 1, wherein  $n$  may be varied based on a rate of change of said content serving sites making up  $n$  content serving sites.

20 12. The method of claim 1, wherein said setup phase occurs when the computer network is first set up.

13. The method of claim 1, wherein said execution phase occurs when the computer network is operating normally.

14. An apparatus for using a phased learning approach for determining the closest

5 content serving sites to a fixed location in a computer network including:

a DNS request receiver coupled to the fixed location;

a response time determiner coupled to said DNS request receiver and coupled to the content serving sites;

10 a query sender coupled to said response time determiner and coupled to the content serving sites;

a data receiver coupled to the fixed location;

a memory coupled to said query sender; and

a data storer coupled to said data receiver and coupled to said memory.

15 15. The apparatus of claim 14, wherein said response time determiner includes an  $n$  fastest content serving site chooser and an  $m$  other content serving site chooser.

16. The apparatus of claim 14, wherein said query sender includes a query timer.

20 17. The apparatus of claim 14, wherein said response time determiner includes a pinger and a reply receiver.

18. The apparatus of claim 14, wherein said query sender includes a response time orderer.

19. The apparatus of claim 14, further including a phase determiner coupled to said response time determiner, said query sender, and said data receiver.

20. A program storage device readable by a machine, tangibly embodying a program of instructions executable by the machine for using a phased learning approach for determining the closest content serving sites to a fixed location in a computer network, the method including:

in a setup phase:

storing data as to the transit times of each of the content serving sites in a table;

in an execution phase:

receiving a DNS request from said fixed location;

determining a response time for each of  $n$  fastest content serving sites and  $m$  other content serving sites, said  $n$  fastest content serving sites chosen by using said data and choosing the  $n$  content serving sites with the lowest transit times, said  $m$  other content serving sites chosen by selecting new content serving sites as well as randomly selecting old content serving sites;

sending queries to each of the  $n$  fastest content serving sites and  $m$  other content serving sites, timing said queries so that they arrive at each of the  $n$  fastest

content serving sites and  $m$  other content serving sites at the same time by using said response time for each of the  $n$  fastest content serving sites and  $m$  other content serving sites;

receiving data from said fixed location as to the transit times of each of the  $n$  fastest content serving sites and  $m$  other content serving sites; and storing said data in said table.

21. An apparatus for using a phased learning approach for determining the closest content serving sites to a fixed location in a computer network including:

means for storing data as to the transit times of each of the content serving sites in a table in a setup phase;

means for receiving a DNS request from said fixed location in an execution phase;

means for determining a response time for each of  $n$  fastest content serving sites and  $m$  other content serving sites, said  $n$  fastest content serving sites chosen by using said data and choosing the  $n$  content serving sites with the lowest transit times, said  $m$  other content serving sites chosen by selecting new content serving sites as well as randomly selecting old content serving sites in an execution phase;

means for sending queries to each of the  $n$  fastest content serving sites and  $m$  other content serving sites, timing said queries so that they arrive at each of the  $n$  fastest content serving sites and  $m$  other content serving sites at the same time by using said response time for each of the  $n$  fastest content serving sites and  $m$  other content serving sites in an execution phase;

means for receiving data from said fixed location as to the transit times of each of the  $n$  fastest content serving sites and  $m$  other content serving sites in an execution phase; and

means for storing said data in said table in an execution phase.

5

22. The apparatus of claim 21, further including:

means for receiving a DNS request from said fixed location;

means for determining a response time for each of the content serving sites;

means for sending queries to each of the content serving sites, timing said queries

so that they arrive at each of the content serving sites at the same time by using said response time for each of the content serving sites;

means for receiving data from said fixed location as to the transit times of each of the content serving sites;

storing said data in a table;

23. The apparatus of claim 22, wherein said means for determining the response time for each of the content serving sites includes:

means for pinging each of the content serving sites;

means for receiving a reply from each of the content serving sites; and

means for determining the response time for each of the content serving sites by computing the difference in time from said pinging to said receiving for each of the content serving sites.

24. The apparatus of claim 21, wherein said means for determining the response time for each of  $n$  fastest content serving sites and  $m$  other content serving sites includes:

means for pinging each of  $n$  fastest content serving sites and  $m$  other content

5 serving sites;

means for receiving a reply from each of  $n$  fastest content serving sites and  $m$  other content serving sites; and

means for determining the response time for each of  $n$  fastest content serving sites and  $m$  other content serving sites by computing the difference in time from said pinging to said receiving for each of the each of said  $n$  fastest content serving sites and  $m$  other content serving sites.

25. The apparatus of claim 21, wherein transit time is the time it takes a content serving site to send a packet to said original location after receiving a query.

26. The apparatus of claim 21, wherein said data includes at least one record, said record including a transit time, a network address for said fixed location, and a network address for a content serving site, said transit time being the transit time between said fixed location and said content serving site.

27. The apparatus of claim 22, wherein said means for sending queries to each of the content serving sites includes:



means for querying each of the content serving sites in order of their response times, longest response time first, setting a delay for a query for a particular content serving site from the time the query to the content serving site with the longest response time as one half of said particular content serving site's response time.

5

28. The apparatus of claim 21, wherein said means for sending queries to each of the  $n$  fastest content serving sites and  $m$  other content serving sites includes:

means for querying each of said content serving sites  $n$  fastest content serving sites and  $m$  other content serving sites in order of their response times, longest response time first, setting a delay for a query for a particular content serving site from the time the query to said  $n$  fastest content serving site or  $m$  other content serving site with the longest response time as one half of said particular content serving site's response time.

29. The apparatus of claim 21, wherein  $n$  is 3 and  $m$  is 2.

30. The apparatus of claim 21, wherein  $n$  may be varied based on stability of the computer network.

31. The apparatus of claim 21, wherein  $n$  may be varied based on a rate of change of

20 said content serving sites making up  $n$  content serving sites.



ABSTRACT

In order to direct content requests to an appropriate content serving site in a computer network, a phased learning approach is utilized to ensure that, as best as possible, the request is made to content serving site with the shortest delay. In a setup phase, an indirect path return geographic sever load balancer times queries to all of the individual content serving sites so that the queries all arrive at the content serving sites at the same time. Therefore, when the requesting fixed location receives communications from the content serving sites, it can easily tell which content serving site has the least delay by an established race condition. The winner of the race may then be relayed to the indirect path return geographic server load balancer for later usage. In an execution mode, only the  $m$  fastest content serving sites and  $n$  other sites (used to test random and new sites) are queried when a DNS request arrives from the requesting fixed location. The particular  $m$  fastest content serving sites and  $n$  other sites may be dynamically updated so as to ensure the most reliable directing of requests. This solution provides a very efficient and effective means by which to determine closest content serving sites while keeping load balancer-created traffic at a minimum.




FIG. 1

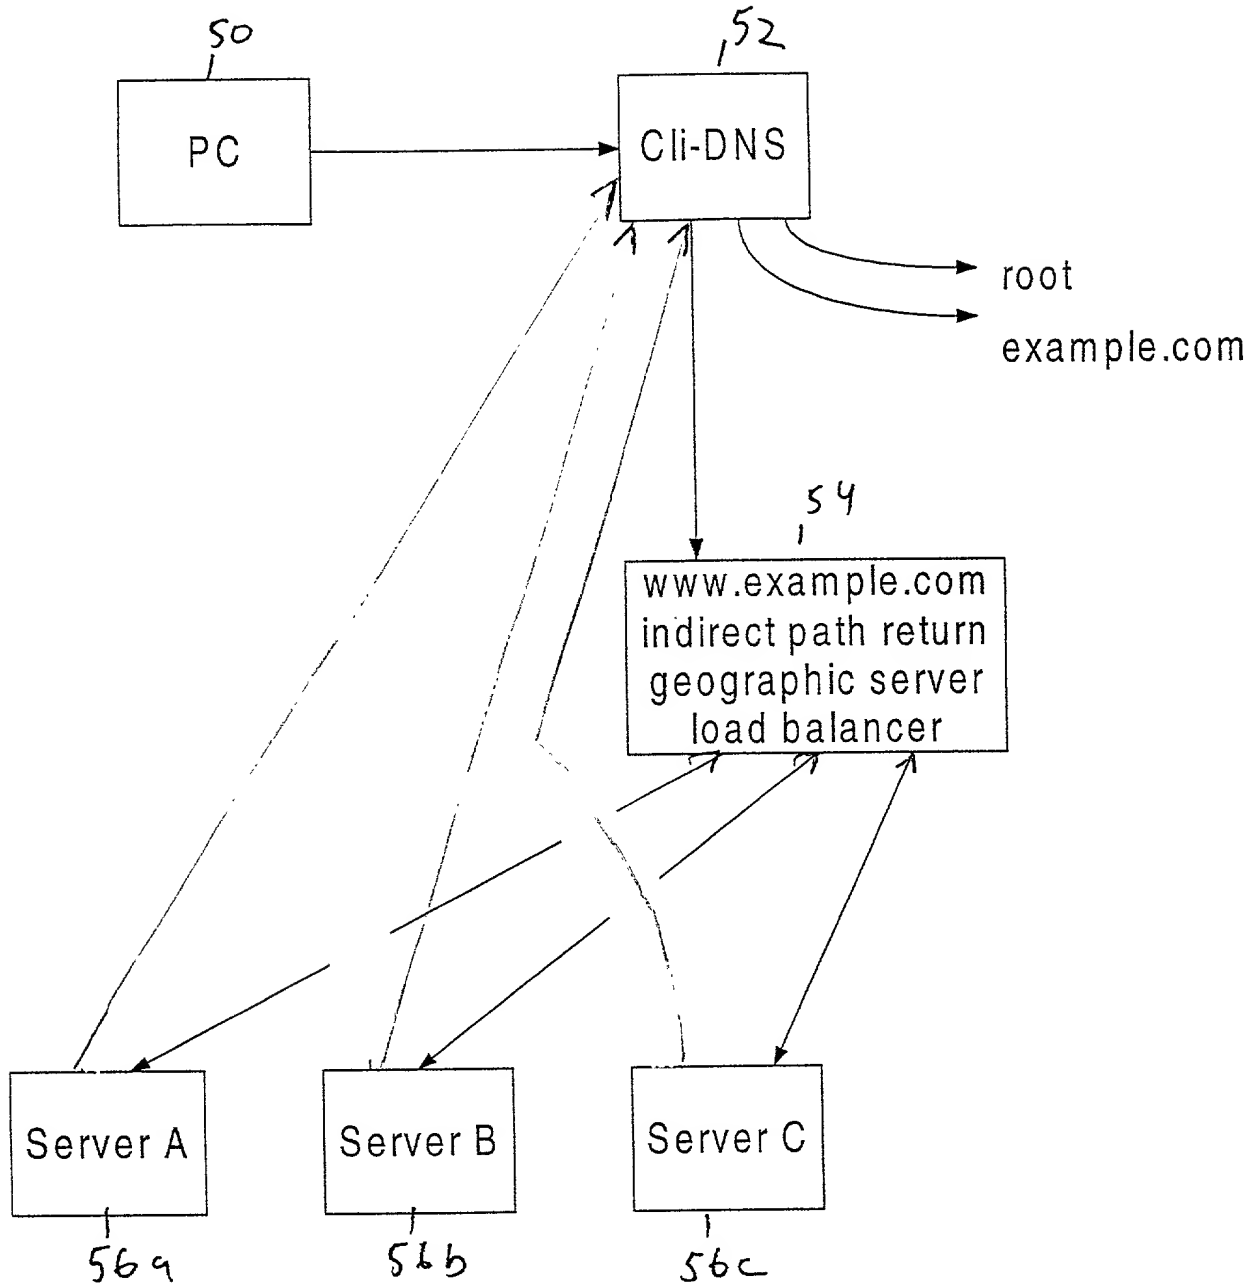
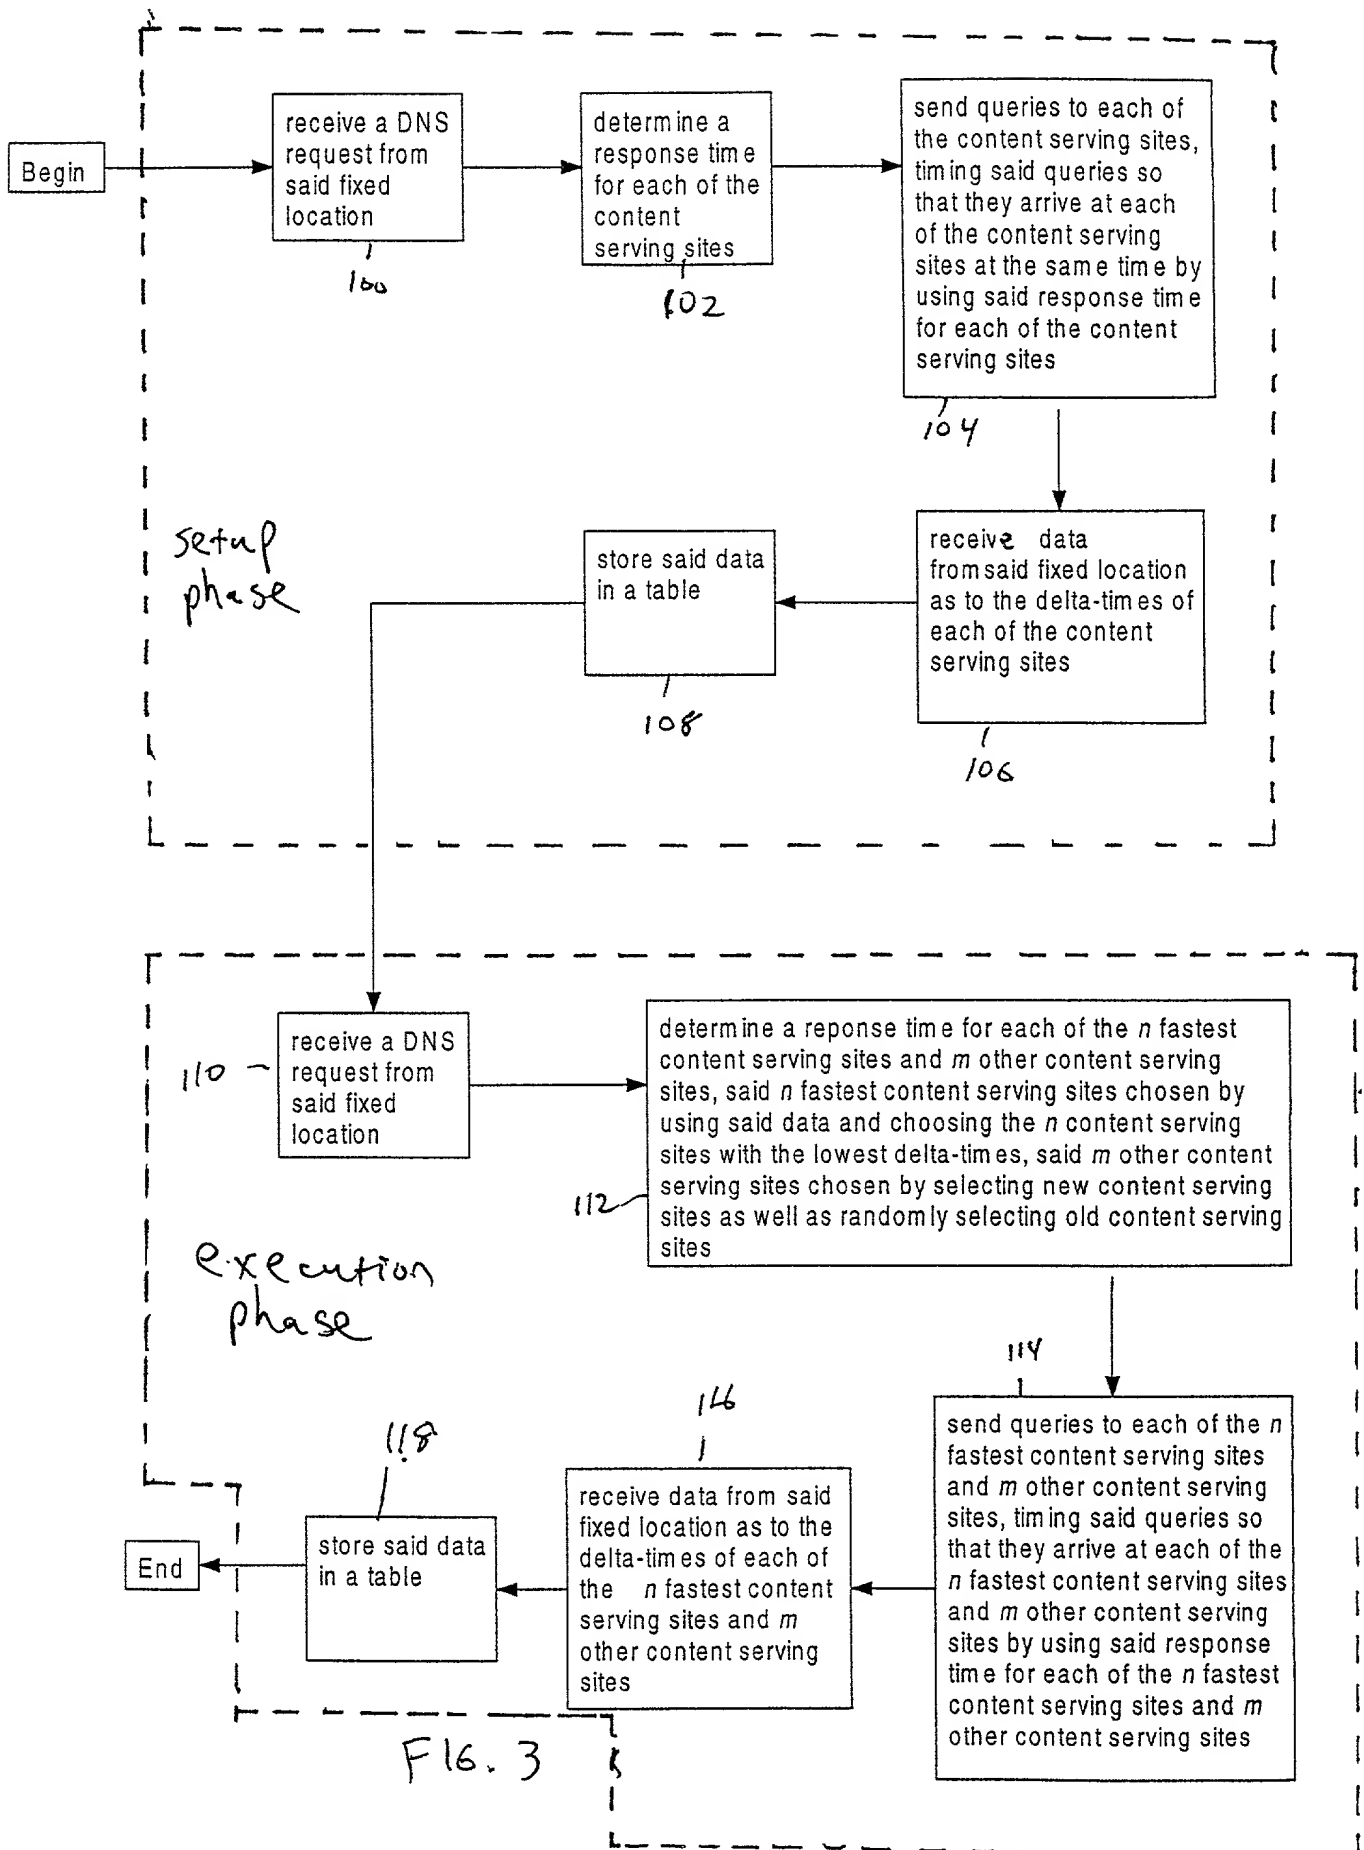
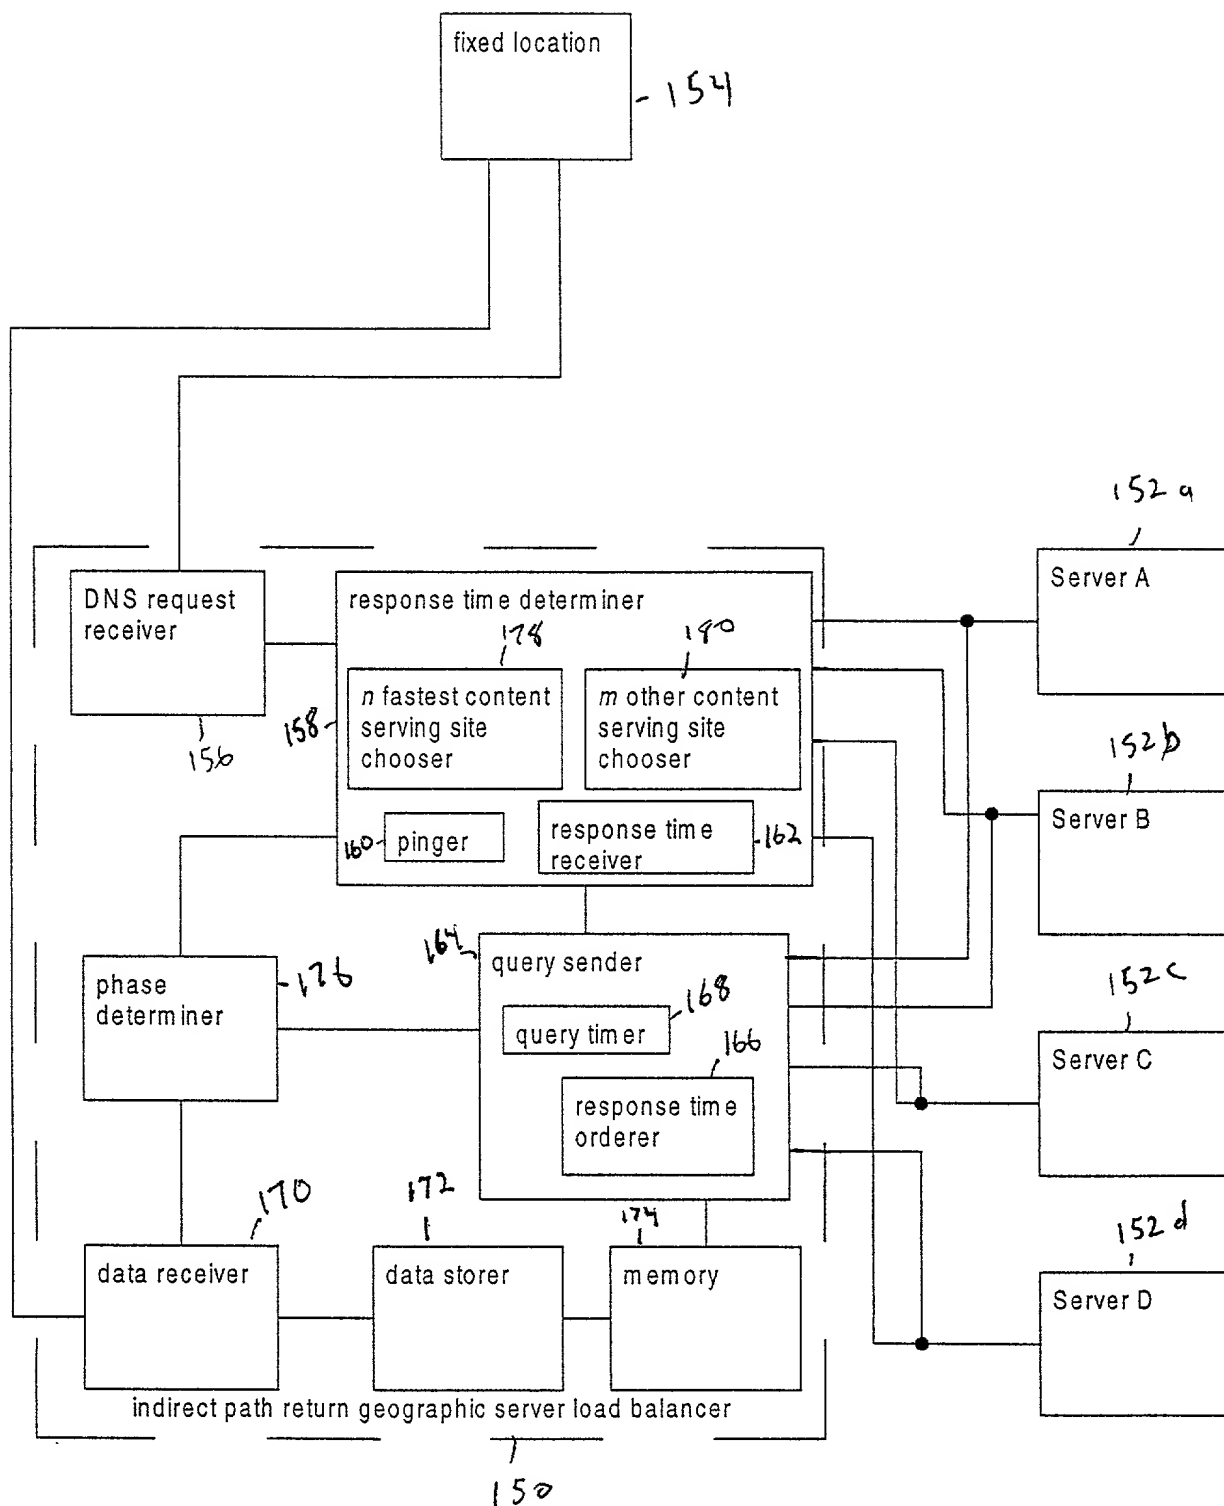


FIG. 2





F16.4